Exact Analysis of TTL Cache Networks. The Case of Caching Policies driven by Stopping Times

Daniel S. Berger
University of Kaiserslautern

Philipp Gland
TU Berlin

Sahil Singla
Carnegie Mellon University

Florin Ciucu
University of Warwick

arXiv:1402.5987v1 [cs.PF] 24 Feb 2014

ABSTRACT

TTL caching models have recently regained significant research interest, largely due to their ability to fit popular caching policies such as LRU. This paper advances the state-of-the-art analysis of TTL-based cache networks by developing two *exact* methods with orthogonal generality and computational complexity. The first method generalizes existing results for line networks under renewal requests to the broad class of caching policies whereby evictions are driven by stopping times. The obtained results are further generalized, using the second method, to feedforward networks with Markov arrival processes (MAP) requests. MAPs are particularly suitable for non-line networks because they are closed not only under superposition and splitting, as known, but also under input-output caching operations as proven herein for phase-type TTL distributions. The crucial benefit of the two closure properties is that they jointly enable the first exact analysis of feedforward networks of TTL caches in great generality.

1. INTRODUCTION

Time-to-Live (TTL) caches decouple the eviction mechanisms amongst objects by associating each object with a timer. When a timer expires the corresponding object is evicted from the cache. This seemingly simple scheme explicitly guarantees (weak) consistency, for which reason it has been widely deployed in DNS and web caching. What has recently however made TTL analytical models quite popular is a subtle mapping between *capacity-driven* (e.g., Least-Recently-Used (LRU)) and *TTL-based* caching policies. This mapping was firstly established through a remarkably accurate approximation by Che *et al.* [12] for the popular LRU policy, which was recently theoretically justified and extended to FIFO (first-in-first-out) and RND (random eviction) policies (Fricker *et al.* [20]), and further confirmed to hold for broader arrival models (Bianchi *et al.* [7]), and even in networks with several replication strategies (Martina *et al.* [30]). Moreover, Fofack *et al.* [19] independently argued that TTL caches capture the properties of LRU, FIFO, and RND policies, and, remarkably, presented the first exact analysis for a line of TTL caches. While the analysis of TTL caches is arguably simpler and more general than the analysis of capacity-driven policies, the exact analysis of TTL networks in general has remained an open problem.

When considering a cache network (including TTL-based), there are two inherent network operations which complicate the analysis: *input-output* and *superposition*. Given a caching node serving a request (point) process for some object (i.e., the *input*), the corresponding miss process is a sample of the request process at those points when the object is absent from the cache. The exact characterization of the *output* process is challenging; for instance, the convenient and often assumed memorylessness property of request processes would not be retained by the corresponding miss process due to the TTL's inherent filtering effect. The *superposition* operation occurs when merging miss processes from upstream caches into a new input process. Since convenient statistical properties of the input processes (e.g., the renewal property) are altered through superposition, the analytical tractability of input-output operations becomes even more complex.

In this paper we give the first exact analysis of caching (feedforward) networks by jointly addressing broad classes of request models, TTL distributions, and caching policies. The request processes are either renewals or Markov arrival processes (MAPs); the latter are dense in a suitable class of point processes and generalize, in particular, the more popular Markov-modulated Poisson processes. The TTLs follow general distributions including phase-type (PH), which are dense within the set of probability distributions on $[0, \infty)$. Moreover, we consider an abstract model for TTL caching policies, whereby cache evictions are driven by stopping times, and which captures in particular three popular policies. The '$\mathcal{R}$' policy regenerates the TTLs at every object's request and maps to the LRU policy. The 'Σ' policy regenerates the TTLs only at those requests resulting in cache misses and maps to FIFO and RND policies. Lastly, the '$\min(\Sigma, \mathcal{R})$' policy combines the key features of '$\mathcal{R}$ ' and 'Σ', i.e., weak consistency guarantee and efficient utilization of cache space, respectively.

We structure our results in two parts. First we generalize the recent results from Fofack *et al.* [19] which cover line networks, renewals requests, general TTL distributions, and the '$\mathcal{R}$' policy, by additionally covering the 'Σ' and '$\min(\Sigma, \mathcal{R})$' policies. The key analytical contribution is a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$10.00.

unified method to recursively characterize input-output operations. This method is based on a suitable change of measure technique using martingales to derive the Laplace transform of a stopped sum, whereby the sum's stopping time characterizes the caching policy. Leveraging certain martingale properties to extend from deterministic to stopping times, we are able to systematically analyze the three caching policies and conceivably many others. The proposed method suffers however from the same limitation as [19]: since renewals are not closed under superposition, unless Poisson, only lines of caches can be (exactly) analyzed.

To address the annoying limitation of renewals' lack of superposition closure, the second part of the paper advocates MAPs to model request processes. The motivation to use MAPs is fairly straightforward since MAPs are known to be closed under superposition. What is remarkable, however, is that we are able to prove that MAPs are also closed under the input-output operation when the TTLs are described by PH distributions, for all the three '$\mathcal{R}$', 'Σ', and '$\min(\Sigma, \mathcal{R})$' policies. In other words, miss processes are also MAPs and trees of cache nodes can be iteratively analyzed. As a side remark, MAPs are also closed under a splitting operation, enabling thus the analysis of feedforward networks.

The two proposed methods advance the state-of-the-art analysis of TTL cache networks by providing the first exact results covering broad request models, caching policies, and network topologies. The second method in particular has the key feature of enabling the first (exact) analysis of feedforward networks and is thus conceivably more general than the first one. However, there is a fundamental trade-off between the two, which is driven by the state explosion of MAPs under the superposition and input-output operations. Therefore, while the first method is computationally fast but restricted to line networks, the second has a much wider applicability but suffers from a high computational complexity, i.e., exponential in the number of caches. However, because the underlying MAPs' matrices are sparse, the numerical complexity can be significantly reduced by using appropriate numerical methods or by formulating lower and upper bound stochastic models.

The rest of the paper is organized as follows. In Section 2 we summarize related cache models and discuss related work. In Section 3 we list model definitions and key objectives for the analysis. In Section 4 we present the change of measure technique to address lines of caches with renewal requests. In Section 5 we consider more general networks with MAP requests. Finally, in Section 6, we conclude the paper. Some of the proofs and examples of MAP input-output constructions are given in the Appendix.

2. CACHE MODELS AND RELATED WORK

Caching is implemented in many computer and communication systems, such as CPUs, databases, or the Internet. Consequently, many analytical models and techniques have been developed to study its performance. Next we discuss related caching policies and results.

Caching policies can roughly be divided into two groups: *capacity-driven* and *TTL-based* (see Rizzo and Vicisano [36]). In the former, objects' evictions are driven by the arrivals of uncached objects and the capacity constraint. In the latter, objects' evictions are determined by individual timers. When compared, TTL-based cache models are typically easier to analyze because the caching behavior of different objects is decoupled and can be thus represented in terms of independent point processes.

In this paper we address the following three TTL-based caching policies, which differ in the behavior of the TTLs' resets and eviction times:

1. *Policy* $\mathcal{R}$: The TTL is reset with every request and an object is evicted upon the TTL's expiration.
2. *Policy* Σ: The TTL is reset only at the times of unsuccessful requests and an object is evicted upon the TTL's expiration.
3. *Policy* $\min(\Sigma, \mathcal{R})$: Two TTLs are reset in parallel according to the $\mathcal{R}$ and Σ policies, respectively, and an object is evicted upon the expiration of *either* of them.

The $\mathcal{R}$ policy can be regarded as a TTL-based correspondent of LRU caches [12]. $\mathcal{R}$ was properly formalized and analyzed by Fofack *et al.* [19]. In turn, the Σ policy can be regarded as a TTL-based correspondent of FIFO and RND (see, e.g., Fricker *et al.* [20]); as a side remark, DNS and web caching implement a variant of Σ.

Unlike $\mathcal{R}$ and Σ, the $\min(\Sigma, \mathcal{R})$ policy has not yet been formalized, although related implementations exist. For instance, the default mechanism in Amazon ElastiCache to enforce an upper bound on memory consumption is to implement an LRU policy on top of evictions caused by TTLs [2]. Another example is the Squid web cache which explicitly uses both $\mathcal{R}$ and Σ TTLs; the former is set locally and referred to as the "Storage LRU Expiration Age" [1], whereas the latter is set by content owners.

In the following we discuss related work starting with single cache models. Early analytical models addressed capacity-driven caches, which, contrary to their implementation simplicity, proved to be difficult to analyze. Some of the classic works (e.g., King [28] and Gelenbe [23]) provided exact results for LRU, FIFO, and RND policies. However, these results were argued to be intractable by Dan and Towsley [16] and Jelenkovic [26], who proposed instead accurate and computationally fast approximations.

With respect to the second group of TTL based caches, the first analytical model for a single Σ cache under renewal arrivals and deterministic TTLs was given by Jung *et al.* [27] who derived the steady-state hit probability. Under the same assumptions, Bahat and Makowski [6] extended this result to the case of non-zero delays between the origin server and the cache, and derived the hit probability for those requests which are consistent with documents that undergo ongoing updates. For a single $\mathcal{R}$ cache, Fofack *et al.* [19] obtained the exact hit probability and other metrics for renewal arrivals.

Next we revisit analytical results for cache networks. Rosenzweig *et al.* [37] applied the approximation scheme from [16] to networks of LRU caches (at the expense however of errors up to 16%). Psaras *et al.* [35] proposed a Markov chain approximation for LRU caches which can then be linked together to form tree networks by assuming each cache's miss process to be Poisson. Under a similar approximation scheme (i.e., each cache's request process is Poisson) Gallo *et al.* [22] considered homogeneous tree networks under the RND policy. A connection between the two domains of capacity-driven and TTL-based policies was established by Che *et al.* [12] for a simple two-level LRU cache network. With the strong case made by Fricker *et al.* [21, 20]

on its wide applicability, the so-called *Che approximation* has recently gained popularity: its impressive accuracy and generality was confirmed by Bianchi *et al.* [7] and Martina *et al.* [30], and parallel extensions to several other caching policies and replication strategies have also been proposed.

The success of the Che approximation is due to a subtle mapping from the domain of capacity-driven caches to the domain of TTL-based caches. In the case of LRU, the key idea is to couple the cache capacity with the durations that objects spent in a cache under the condition that no further arrivals occur (i.e., $\mathcal{R}$). By assuming these (random) durations as deterministic and equal for every object, the LRU model reduces to a TTL model [12, 20] that is easier to analyze.

To analyze TTL cache networks (e.g., as arising from the Che approximation for a capacity-driven policy), Martina *et al.* [30] rely on Poisson approximations of the output processes (for both $\mathcal{R}$ and Σ caches) and report accurate results. Moreover, Fofack *et al.* [19] derived the first exact results for a line of $\mathcal{R}$ caches and analyzed tree networks by relying on a renewal approximation of superposed processes.

Unlike these works, which assume an independent cache behavior, another set of works considered hierarchies of caches where the layers are synchronized by an aging mechanism: the TTL value at child caches are set to coincide with the remaining TTL of parent caches. In this way, Cohen and Kaplan [14] were able to derive the miss rate for a two-level hierarchy, and Cohen *et al.* [13] extended this result to heterogeneous parent nodes. Remarkably, by ingeniously leveraging the system's property that misses occur synchronously, Hou *et al.* [25] were able to analyze trees of caches under Poisson arrivals at the leaf caches.

We finally overview some recent related work on potential applications. Borst *et al.* [9] addressed the optimization of link utilization with a linear program for content placement in a tree network. An important insight is that placing caches close to the network edge often supersedes placing additional caches within the network. This was further confirmed through simulations by Psaras *et al.* [35] and Fayazbakhsh *et al.* [18] for LRU caches; in particular, the latter argued that placing caches at the leaf nodes of an ISP access tree does not sacrifice performance when compared to omnipresent cache placement. A potentially interesting application of our results is to analytically confirm whether this insight holds in a much broader sense, e.g., for other network topologies where different nodes can implement different caching policies (i.e., either $\mathcal{R}$, Σ, or $\min(\Sigma, \mathcal{R})$).

3. ROADMAP

In this section we state the key objectives for analyzing lines of caches and feedforward networks. First we give some general definitions concerning some arbitrary node in a cache network.

DEFINITION 1 (ARRIVAL/INPUT PROCESS).
For each object, the arrival process is represented as a point (counting) process $N(t)$. The corresponding inter-arrival process is denoted by $\{X_t\}_{t\geq 1}$.

When analyzing a cache network, one needs to characterize the miss/output process relating two consecutive caches. We give the definition in the renewal case.

DEFINITION 2 (MISS/OUTPUT PROCESS).
Let a caching node with inter-request times and TTL's be given by two independent renewal processes: $\{X_t\}_{t\geq 1}$ and $\{T_t\}_{t\geq 1}$. The corresponding miss process is also a renewal process with the same distribution as the stopped sum

$$S_\tau := X_1 + \cdots + X_\tau \ , \tag{1}$$

where τ is a stopping time defined separately for each caching (TTL) policy:

$$\textit{Policy } \mathcal{R}: \qquad \tau := \min\{t : X_t > T_t\} \tag{2}$$

$$\textit{Policy } \Sigma: \qquad \tau := \min\{t : \sum_{s=1}^{t} X_s > T_1\} \tag{3}$$

$$\textit{Policy } \min(\Sigma,\mathcal{R}): \qquad \tau := \min\{\min\{t : \sum_{s=1}^{t} X_s > T_1^{\Sigma}\}, \min\{t : X_t > T_t^{\mathcal{R}}\}\} \ . \tag{4}$$

For the last policy, T_1^{Σ} and $T_t^{\mathcal{R}}$ are independent renewal processes.

The corresponding definition for the non-renewal case can be stated similarly and is omitted for brevity.

As a side remark, the structure of the first two stopping times justifies the notation for the caching policies by $\mathcal{R}$ and Σ: the former has a renewal structure (whence the letter $\mathcal{R}$), whereas the latter has a sum structure (whence the letter Σ). We incorporate this notation in the whole notation of a caching node, by borrowing from Kendall's notation in queueing theory.

NOTATION 1 (CACHING NODE).
Depending on the TTL policy, a caching node is denoted as either one of the triplets

$$G\text{-}G\text{-}\mathcal{R} \textit{ or } G\text{-}G\text{-}\Sigma \textit{ or } G\text{-}G\text{-}\textit{min}(\mathcal{R},\Sigma) \ ,$$

where the two G's stand for the generic distributions of the inter-arrival times and the TTLs, respectively.

As an example, a cache with exponentially distributed inter-arrival and TTL times, and implementing the $\mathcal{R}$ policy, is denoted by $M\text{-}M\text{-}\mathcal{R}$. Some other distributions used in this paper are the deterministic (D) case, the exponential (M), and the phase-type (PH) distribution.

Cache performance is commonly measured in terms of hit/miss probabilities, which indicate the improvement in link utilization when using a cache. For example, the cost metrics in [18] depend on the hit/miss probabilities and the arrival rates at the leaves of a caching tree.

DEFINITION 3 (HIT/MISS PROBABILITY).
Consider an arbitrary cache with arrival process $N(t)$ and miss process $M(t)$, for some fixed object. The hit and miss probabilities are defined as

$$H := \lim_{t\to\infty} \left(1 - \frac{M(t)}{N(t)}\right)$$

$$M := \lim_{t\to\infty} \frac{M(t)}{N(t)} \ ,$$

respectively, subject to convergence.

Another metric of interest is the cache occupancy which defines the average amount of storage required by an object,

and also establishes a connection between capacity-driven and TTL-based cache models through a suitable set of parameters. For example, the connection between an LRU cache's capacity and the corresponding TTL-$\mathcal{R}$ model is established by equalizing the summation of the occupancies of the objects in the TTL cache with the LRU capacity [12].

DEFINITION 4 (CACHE OCCUPANCY).
Let $C(t)$ be a random binary process representing whether the object is in the cache or not at time t. The cache occupancy is defined as

$$\pi := \lim_{t\to\infty} \frac{\int_0^t C(s)ds}{t} \ .$$

We next briefly introduce the key objectives for analyzing lines and feedforward networks.

3.1 Lines of Caches with Renewal Arrivals

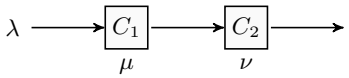

Figure 1: A line of two caches C_1 and C_2 with Poisson arrivals with rate λ and exponential TTLs with rates μ and ν

Consider the simplified line network with two nodes from Figure 1. At the first node, requests for some object arrive according to a renewal process $\{X_t\}_{t\geq 1}$. If the object is in the cache at the time of a request, then the request is successful and the object is fetched. Otherwise, for every unsuccessful request at some node, the object is recursively requested at the next node in the line. Once the object is successfully found at some downstream node, it is (instantaneously) transferred to the upstream nodes. For the model's completeness we assume that the last node always has a copy of the object. All the nodes implement one of the $\mathcal{R}$, Σ, or $\min(\Sigma, \mathcal{R})$ caching policies (different nodes are allowed to implement different policies).

We address a single node with a given arrival/input process and TTL distribution. For this setting our objective is to derive the Laplace transform of the corresponding miss/output process, i.e., of the stopped random walk S_τ from Eq. (1):

$$\mathcal{L}_\omega(S_\tau) := E\left[e^{-\omega(X_1+X_2+\cdots+X_\tau)}\right] \ . \tag{5}$$

This technique can be iteratively applied along an entire line of caches using numerical methods (as in Fofack *et al.* [19]); note that numerical methods are not necessary in the case of exponential TTLs. We emphasize that, unlike [19] which is restricted to $\mathcal{R}$, our method additionally covers Σ and $\min(\Sigma, \mathcal{R})$.

3.2 Feedforward Networks with MAP Arrivals

The previous technique suffers from the major limitation that cache requests must be a renewal process and thus it does not apply to more general topologies subject to a superposition operation.

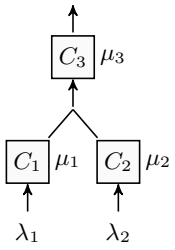

Figure 2: A tree of caches

Indeed, consider the simple tree topology from Figure 2 in which the inter-request times at the leaf caches are $\sim$ exp(λ_1) (exponential) and $\sim$ exp(λ_2), and the corresponding TTLs are $\sim$ exp(μ_1) and $\sim$ exp(μ_2), respectively; all the processes are independent. The inter-miss times at the leaf caches are renewal processes with $hypo(\lambda_1, \mu_1)$ (hypoexponential) and $hypo(\lambda_2, \mu_2)$ distributions (to be shown in Table 1). The superposition of the two renewal processes is not a renewal process, which means that the technique targeting lines of caches does not apply at the root cache (the superposition of renewal processes is a renewal process if and only if the superposed processes are Poisson).

Let us now recall the two main analytical operations which must be accounted for when analyzing trees of caches:

1. *input-output*: the characterization of the inter-miss process from the inter-request process.
2. *superposition*: the characterization of a single inter-request process from multiple ones (e.g., at the root cache from Figure 2).

Unlike the *input-output* operation which will be shown in Section 4 to be tractable (yet subject to recursions and also evaluation of convolutions in the case of Σ and $\min(\Sigma, \mathcal{R})$ caches), the *superposition* operation is conceivably the bottleneck due to the lack of closure for renewal processes. To circumvent this apparent difficulty, the natural generalization of renewal processes are Markov arrival process (MAPs), which are known to be closed under *superposition* (and also splitting). The remaining objective is to additionally show that MAPs are also closed under the *input-output* operation of caches (see Section 5).

We point out that the apparently straightforward idea of using MAPs has been efficiently used in the past to model systems with non-renewal behavior, e.g., single queues with non-renewal arrivals (Lucantoni *et al.* [29]) or closed queueing networks with non-renewal workloads (Casale *et al.* [11]); for an excellent related survey see Asmussen [3].

4. LINES OF CACHES

In this section we propose a unified method to analyze lines of *G-G-*$\mathcal{R}$, *G-G-*Σ, and *G-G-*$\min(\Sigma, \mathcal{R})$ caches. First we instantiate the caching metrics from Definitions 3 and 4 for the renewal case.

LEMMA 1 (HIT/MISS PROBABILITIES).
For G-G-$\mathcal{R}$*, G-G-*Σ*, and G-G-min*$(\mathcal{R}, \Sigma)$ *caches, the hit and miss probabilities from Definition 3 become*

$$H = \frac{\mathbb{E}[\tau] - 1}{\mathbb{E}[\tau]} \ \textit{and}\ M = \frac{1}{\mathbb{E}[\tau]} \ . \tag{6}$$

In particular, for G-G-$\mathcal{R}$*, it holds*

$$H = \mathbb{P}(X \leq T) \ \textit{and}\ M = \mathbb{P}(X > T) \ . \tag{7}$$

The expression of the miss probability for the *G-G-*Σ cache is the same as in Jung *et al.* [27]. Unlike in the *G-G-*$\mathcal{R}$ case, $E[\tau]$ for the other two policies cannot be generally given in closed-form due to the underlying convolution in

the definition of τ from Eqs. (3)-(4); it is, however, often straightforward to derive $E[\tau]$ for particular distributions of $\{X_t\}_{t\geq 1}$ and $\{T_t\}_{t\geq 1}$.

PROOF. Let t_n denote the point process of the unsuccessful request times (i.e., the miss times) for $n \geq 1$ and $t_0 = 0$. Using the renewal property of $\{X_t\}_{t\geq 1}$ and $\{T_t\}_{t\geq 1}$, and the strong law of large numbers, we have

$$\lim_{n\to\infty} \frac{M(t_n)}{N(t_n)} = \lim_{n\to\infty} \frac{n}{\tau_1 + \cdots + \tau_n} = \frac{1}{E[\tau]} ,$$

where τ_i denotes the stationary sequence of stopping times, as defined in Eqs. (2) and (3), but starting from $t \geq t_{i-1}$ in the usual renewal sense. Moreover, since for $t \in (t_{i-1}, t_i]$

$$\frac{M(t_i) - 1}{N(t_i)} < \frac{M(t)}{N(t)} \leq \frac{M(t_{i-1}) + 1}{N(t_{i-1})} ,$$

the limit $\lim_{t\to\infty} \frac{M(t)}{N(t)}$ exists and Eq. (6) is proven. The particular expression for G-G-$\mathcal{R}$ cache follows directly from the geometric distribution of τ. □

LEMMA 2 (CACHE OCCUPANCY).
For the G-G-$\mathcal{R}$, G-G-Σ, and G-G-$\min(\Sigma, \mathcal{R})$ caches, the cache occupancies from Definition 4 become

$$\pi_{\mathcal{R}} = \frac{\mathbb{E}[min\{X, T\}]}{\mathbb{E}[X]},\ \pi_{\Sigma} = \frac{\mathbb{E}[T]}{\mathbb{E}[S_\tau]}$$

$$\pi_{\min(\Sigma,\mathcal{R})} = \frac{E\left[\min\{\sum_{s=1}^{\tau} \min\{X_s, T_s^{\mathcal{R}}\}, T_1^{\Sigma}\}\right]}{E[S_\tau]} .$$

The expression for G-G-$\mathcal{R}$ is the same as the one given by Fofack *et al.* [19] (written therein in the equivalent form $\pi_{\mathcal{R}} = \mathbb{E}\left[\int_0^X \mathbb{P}(T > t)dt\right] / \mathbb{E}[X]$). Note that the last expectation depends on the stopping time τ from Eq. (4), whose mass function is later provided in Corollary 3; moreover, to compute the cache occupancy, a decoupling argument like the one we provide for the transforms of the inter-miss times is needed to avoid the implicit correlations amongst the stopping time, the inter-arrival times, and the TTLs (for all $\mathcal{R}$, Σ, and $\min(\Sigma, \mathcal{R})$).

PROOF. In the case of G-G-$\mathcal{R}$, denote the point process t_n of the request times, i.e., $t_n = \sum_{i=1}^{n} X_i$. Using the renewal property of $\{X_t\}_{t\geq 1}$ and $\{T_t\}_{t\geq 1}$, and the strong law of large numbers, we have

$$\lim_{n\to\infty} \frac{\int_0^{t_n} C(s)ds}{t_n} = \lim_{n\to\infty} \frac{\sum_{i=1}^{n} \min\{X_i, T_i\}}{\sum_{i=1}^{n} X_i} = \frac{E[\min\{X, T\}]}{E[X]} .$$

In the case of G-G-Σ, we use the same embedding t_n as in the proof of Lemma 1 such that

$$\lim_{n\to\infty} \frac{\int_0^{t_n} C(s)ds}{t_n} = \lim_{n\to\infty} \frac{\sum_{i=1}^{n} T_i}{\sum_{i=1}^{n} S_{\tau_i}} = \frac{E[T]}{E[S_\tau]} ,$$

where $S_{\tau_i} := X_{\tau_{i-1}+1} + \cdots + X_{\tau_{i-1}+\tau_i}$.

In both cases, the extensions of the limits to the whole line follows by a bounding argument as in the proof of Lemma 1. The proof for G-G-$\min(\Sigma, \mathcal{R})$ follows using the same embedding points as for G-G-Σ. □

The key problem to obtain caching metrics at the downstream nodes in a line of cache is to recursively characterize the inter-miss time S_τ defined in Eq. (1). The expression of S_τ, as well as those of the above cache metrics, suggests following a martingale based technique to characterize S_τ. This (rough) idea is driven by the fact that stopping times—which are at the core of the very definition of S_τ—preserve certain martingale results, e.g., if L_t is a martingale and τ is a bounded stopping time then $E[L_\tau] = E[L_1]$, which is a particular case of the optional stopping theorem. Note however that caching stopping times are not necessarily bounded, and thus require a more general framework. For example, the stopping times from Eqs. (2)-(4) are under realistic assumptions almost surely finite but may be unbounded.

Next we will demonstrate the effectiveness of relying on martingale techniques to derive an elegant and unified analysis of G-G-$\mathcal{R}$, G-G-Σ, and G-G-$\min(\Sigma, \mathcal{R})$ caches. To this end, we first provide a closed-form result (in most of the cases) for the Laplace transform of the stopped random sum S_τ. This result will be instrumental to the analysis of the $\mathcal{R}$, Σ, and $\min(\Sigma, \mathcal{R})$ policies.

4.1 The Laplace Transform of a Stopped Sum

Consider the two independent renewal processes $\{X_t\}_{t\geq 1}$ and $\{T_t\}_{t\geq 1}$ on a joint probability space $(\Omega, \mathcal{F}, \mathbb{P})$ (e.g., as in Definition 2). Denote the corresponding distribution functions by $F(x)$ and $G(x)$, and assume the existence of corresponding densities $f(x)$ and $g(x)$, respectively. Let $\mathcal{F}_t = \sigma((X_1, T_1), \ldots, (X_t, T_t))$, $\mathbb{F} = (\mathcal{F}, \{\mathcal{F}_t\}_{t\geq 1})$ denote the filtration associated with S_τ, and let $(\Omega, \mathbb{F}, \mathbb{P})$ denote the corresponding filtered probability space. When clear from the context the time indexes are suppressed.

Next we provide a closed-form expression for the Laplace transform of the stopped random walk S_τ from Eq. (5). Recall from Eqs. (2)-(4) that τ is a stopping time with respect to the filtration $\mathcal{F}_t$. One may remark that, due to the intrinsic dependencies amongst X_t's and the stopping time τ, the analysis of the stopped sum S_τ is conceivably quite involved even for stopping times w.r.t. the filtration $\mathcal{F}'_t = \sigma(X_1, X_2, \ldots, X_t)$. In fact, unlike the first moment which is relatively easily obtained as Wald's equation, i.e., $E[S_\tau] = E[\tau]E[X_1]$ (under the additional condition that $E[\tau] < \infty$), higher moments, however, are typically only known in terms of bounds (see Gut [24], p. 22).

Despite the apparent technical difficulties, we will next show that $\mathcal{L}_\omega(S_\tau)$, for stopping times τ w.r.t. $\mathcal{F}_t$, can be derived in a rather straightforward manner. The key idea is to construct a suitable new filtered probability space $(\Omega, \mathbb{F}, \tilde{\mathbb{P}})$, whereby the new probability measure $\tilde{\mathbb{P}}$ decouples the dependencies amongst X_t's and τ. Informally, the key idea to obtain $\mathcal{L}_\omega(S_\tau)$ in closed-form is to *offshore* the underlying derivations into the new $(\Omega, \mathbb{F}, \tilde{\mathbb{P}})$ space.

This technique is known as *change of measure*. The change of measure itself (e.g., from $\mathbb{P}$ to $\tilde{\mathbb{P}}$) is performed as such *measures* in the original space (e.g., $\mathbb{P}(A)$ for $A \in \mathcal{F}$) can be obtained in terms of the new (changed) measure in a much simpler manner. An example of an application of this technique is in rare events simulations, whereby rare events become more likely to occur under the new (changed) measure, or more precisely under the new (changed) density, guaranteeing thus faster convergence speeds than Monte-Carlo simulations (see Pham [34]). Another application is

in pricing risks in incomplete markets, by constructing a new risk-neutral probability measure (see Cox *et al.* [15]). Such risk-neutral measures have also been constructed in financial models, in order to simplify a model with drift into one with constant expectation and allowing thus the application of the Girsanov theorem to describe the process dynamics (see Musiela and Rutkowski [31]). Another application is an elegant proof for Cramér's theorem in large deviation theory (see Dembo and Zeitouni [17], p. 27).

To perform the intended change of measure, we extend the measure construction for a filtration $\mathcal{F}'_t = \sigma(X_1, X_2, \ldots, X_t)$ (see Asmussen [4], p. 358) to the product filtration $\mathcal{F}_t = \sigma((X_1, T_1), \ldots, (X_t, T_t))$ which accounts for $(T_t)_{t\geq 1}$ as well. While the extension proceeds mutatis mutandis, mainly due to the independence between $(X_t)_{t\geq 1}$ and $(T_t)_{t\geq 1}$, the key to our construction is to only *tilt* the distribution $F(x)$ of X_t while preserving the distribution $G(x)$ of T_t; for this reason, we refer to our change of measure as a *fractional* change of measure.

Definition 5 (Fractional Change of Measure). *For any $F \in \mathcal{F}_t$ define the tilted probability measure $\tilde{\mathbb{P}}_t$ as*

$$\tilde{\mathbb{P}}_t(F) := E\left[L_t 1_F\right] \ , \tag{8}$$

where L_t is the Wald's martingale

$$L_t := \frac{e^{-\omega S_t}}{\mathcal{L}_\omega(X)^t} \ \forall \ t \geq 1 \tag{9}$$

w.r.t. the original filtered space $(\Omega, (\mathcal{F}, \{\mathcal{F}'_t\}_{t\geq 1}), \mathbb{P})$, and for some fixed $\omega \in \mathbb{R}$.

The tilted measures $\tilde{\mathbb{P}}_t$, which are by construction restricted to $\mathcal{F}_t$, uniquely extend to a probability measure $\tilde{\mathbb{P}}$ on $\mathcal{F}$ which is Kolmogorov consistent, i.e.,

$$\tilde{\mathbb{P}}(F) = \tilde{\mathbb{P}}_t(F) = E\left[L_t 1_F\right] \ ,$$

for all $F \in \mathcal{F}_t$. The proof follows the proof of Proposition 3.1 from [4], with the observation that L_t is also a martingale w.r.t. the product filtered space $(\Omega, \mathbb{F}, \mathbb{P})$ due to the independence between $(X_t)_{t\geq 1}$ and $(T_t)_{t\geq 1}$.

Besides enabling the construction of the consistent probability measure $\tilde{\mathbb{P}}$ (mainly using the fact that L_t's are martingales with $E[L_t] = 1$), there are two technical reasons behind the fractional change of measure from Definition 5. On one hand, L_t corresponds to the Radon-Nikodym density of the Kolmogorov extended measure $\tilde{\mathbb{P}}$ (in addition to that of $\mathbb{P}_t$ as well) w.r.t. $\mathbb{P}$, i.e., $L_t = \frac{d\tilde{\mathbb{P}}}{d\mathbb{P}}$ on $(\Omega, \mathcal{F}_t)$ for all $t \geq 1$. This allows the computation of integrals w.r.t. $\tilde{\mathbb{P}}$ according to the integration rule

$$\int_A Y d\tilde{\mathbb{P}} = \int_A Y L_t d\mathbb{P} \ \forall A \in \mathcal{F}_t$$

for $\mathcal{F}_t$-measurable Y, under the condition that YL_t is integrable w.r.t. $\mathbb{P}$ (see Billingsley [8], Theorem 16.11). In particular, one has in terms of expectations

$$\tilde{E}[Y] = E[YL_t] \ , \tag{10}$$

where $\tilde{E}[\cdot]$ is the expectation w.r.t. $\tilde{\mathbb{P}}$.

On the other hand, the *particular* expression of L_t from Eq. (9) lends itself, by plugging in above $Y := \mathcal{L}_\omega(X)^t$ and cancelling out terms, to the sought result, i.e.,

$$E\left[e^{-\omega S_t}\right] = \tilde{E}\left[\mathcal{L}_\omega(X)^t\right] \ \forall t \geq 1 \ .$$

The final result, i.e., when t is replaced by a stopping time τ, follows similarly by applying Theorem 3.2 from Asmussen [4].

Theorem 1 (Laplace Transform of S_τ). *For an (a.s.) finite stopping time τ, the Laplace transform of the stopped sum S_τ from Eq. (1) is given by*

$$\mathcal{L}_\omega(S_\tau) = \tilde{E}\left[\mathcal{L}_\omega(X)^\tau\right] \ . \tag{11}$$

Note that this result is a manifestation of the earlier stated motivation that 'stopping times preserve martingale properties', justifying thus our overall martingale framework to analyze the inter-miss times S_τ.

Proof. [1] Fix $T \geq 0$ and choose $Y := \mathcal{L}_\omega(X)^\tau 1_{\{\tau\leq T\}}$ which is $\mathcal{F}_T$ measurable. Applying the integration rule from Eq. (10) and the properties of conditional expectation we get

$$\begin{aligned}
\tilde{E}\left[\mathcal{L}_\omega(X)^\tau 1_{\{\tau\leq T\}}\right] &= E\left[\mathcal{L}_\omega(X)^\tau 1_{\{\tau\leq T\}} L_T\right] \\
&= E\left[E\left[\mathcal{L}_\omega(X)^\tau 1_{\{\tau\leq T\}} L_T \mid \mathcal{F}_\tau\right]\right] \\
&= E\left[\mathcal{L}_\omega(X)^\tau 1_{\{\tau\leq T\}} E\left[L_T \mid \mathcal{F}_\tau\right]\right] \\
&= E\left[e^{-\omega S_\tau} 1_{\{\tau\leq T\}}\right] \ .
\end{aligned}$$

In the last line we used the martingale property of L_T, i.e., $E[L_T \mid \mathcal{F}_\tau] = L_\tau$. From the monotonicity of $1_{\{\tau\leq T\}}$ in T, the proof is complete by applying Lebesgue's dominated convergence theorem (see Theorem 16.4 in Billingsley [8]). □

The crucial aspect about Eq. (11) is that $\mathcal{L}_\omega(X)$ is computed w.r.t. the original probability measure $\mathbb{P}$. What remains to compute is τ's pmf under the changed measure $\tilde{\mathbb{P}}$. In other words, the computations for $\mathcal{L}_\omega(X)$ and the pmf of τ under $\tilde{\mathbb{P}}$ are entirely decoupled, circumventing thus the dependencies in the stopped sum S_τ.

To facilitate the auxiliary calculus under $\tilde{\mathbb{P}}$ we next give the following technical result whose proof is deferred to Appendix A.1.

Proposition 1. *On the new probability space $(\Omega, \mathcal{F}, \tilde{\mathbb{P}})$, the random variables X_t and T_t have the following distribution functions for all $t \geq 1$ and $x \geq 0$*

$$\begin{aligned}
\tilde{F}(x) &:= \tilde{\mathbb{P}}\left(X_t \leq x\right) = \frac{E\left[e^{-\omega X_t} 1_{\{X_t\leq x\}}\right]}{\mathcal{L}_\omega(X)} \\
\tilde{G}(x) &:= \tilde{\mathbb{P}}\left(T_t \leq x\right) = G(x) \ .
\end{aligned}$$

The corresponding densities are $\tilde{f}(x) := d\tilde{F}(x) = \frac{e^{-\omega x} f(x)}{\mathcal{L}_\omega(X)}$ and $\tilde{g}(x) := d\tilde{G}(x) = g(x)$, respectively. Moreover, X_t and T_t remain independent under $\tilde{\mathbb{P}}$.

4.2 The Laplace Transform of Inter-Miss Times

Here we apply Theorem 1 to derive the particular transforms of the inter-miss times for the G-G-$\mathcal{R}$ and G-G-Σ caching model. As the result for the G-G-$\min(\mathcal{R}, \Sigma)$ cache model is notationally complex, it is stated in Appendix A.4. Note that the stopped sum S_τ from Eq. (1) corresponds to the inter-miss time at a G-G-$\mathcal{R}$, G-G-Σ, or G-G-$\min(\Sigma, \mathcal{R})$ cache, depending whether the stopping time τ is defined as in Eqs. (2)-(4), respectively.

[1] The proof proceeds along the same lines as in [4] with the main difference of working on an extended product filtration. Theorem 1 herein is thus a simple extension of Theorem 3.2 from [4], and no significant technical contribution is accordingly claimed.

	M-M-$\mathcal{R}$ = M-M-Σ	M-D-$\mathcal{R}$	M-D-Σ	M-M-$\min(\Sigma,\mathcal{R})$
$\mathcal{L}_\omega(S_\tau)$	$\frac{\lambda}{\lambda+\omega}\frac{\mu}{\mu+\omega}$	$\frac{\lambda e^{-\lambda/\mu}}{\lambda e^{-\lambda/\mu}+\omega e^{\omega/\mu}}$	$\frac{\lambda}{\lambda+\omega}e^{-\omega/\mu}$	$\frac{\lambda}{\lambda+\omega}\frac{\mu+\nu}{\mu+\nu+\omega}$
$\mathbb{P}_{S_\tau}$	$\mathbb{P}_{X_1+T_1}$ $(hypo(\lambda,\mu))$	open	$\mathbb{P}_{X_1+1/\mu}$	$hypo(\lambda,\mu+\nu)$
$\mathbb{E}[S_\tau]$	$\frac{\lambda+\mu}{\lambda\mu}$	$\frac{1}{\lambda e^{-\lambda/\mu}}$	$\frac{\lambda+\mu}{\lambda\mu}$	$\frac{\lambda+\mu+\nu}{\lambda(\mu+\nu)}$
$\pi=h_p$	$\frac{\lambda}{\lambda+\mu}$	$1-e^{-\lambda/\mu}$	$\frac{\lambda}{\lambda+\mu}$	$\frac{\lambda}{\lambda+\mu+\nu}$

Table 1: Explicit Laplace transforms $\mathcal{L}_\omega(S_\tau)$, distribution law $\mathbb{P}$, expectation of the stopped sum S_τ, and cache occupancy for several caching models. Corresponding results for M-D-$min(\Sigma,\mathcal{R})$ are omitted as they depend on a closed form of $\mathbb{P}_{S_\tau}$ for M-D-$\mathcal{R}$ which currently remains open. ($E[X_1]=1/\lambda$, deterministic $T=E[T_1]=E[T_1^\Sigma]=1/\mu$, $E[T_1^\mathcal{R}]=1/\nu$)

COROLLARY 1 (G-G-$\mathcal{R}$). [2]
Let τ as in Eq. (2). If $\psi(\omega) := \mathbb{E}\left[e^{-\omega X}1_{\{X\le T\}}\right] < 1$ for some $\omega > 0$, then the Laplace transform of the inter-miss time in the G-G-$\mathcal{R}$ model is given by

$$\mathbb{E}\left[e^{-\omega S_\tau}\right] = \frac{\mathcal{L}_\omega(X)-\psi(\omega)}{1-\psi(\omega)}\ . \tag{12}$$

To derive Corollary 1, it is sufficient to derive the probability mass function of τ under $\tilde{\mathbb{P}}$, which follows a simple geometric structure. A full proof is stated in Appendix A.2.

COROLLARY 2 (G-G-Σ).
Let τ as in Eq. (3). Then for some $\omega > 0$ the Laplace transform of the inter-miss time in the G-G-Σ model is given by

$$\mathbb{E}\left[e^{-\omega S_\tau}\right] = \sum_{t\ge 1}\phi(\omega)^t\mathbb{E}\left[\tilde{F}^{t-1}(T)\ -\ \tilde{F}^{t}(T)\right] \tag{13}$$

where $\tilde{F}^t$ is the distribution of the t-fold convolution of X in the tilded probability space.

Unlike the G-G-$\mathcal{R}$ model, the G-G-Σ model is more tedious to analyze due to the expression of the stopping time τ from Eq. (3). In particular, to account for the sum in the expression of τ a convolution density is required. We state an according definition and a subsequent proof of Corollary 2 in Appendix A.3.

4.3 Examples

Here we instantiate the results from Corollaries 1, 2, and 3 (the latter is given in Section A.4). In particular, Table 1 gives explicit expressions for the Laplace transforms, distribution laws $\mathbb{P}$, and expectations of the stopped sum S_τ, and also the cache occupancies for several simple caching models. Recall that the results for the $\mathcal{R}$ model have been previously obtained by Fofack *et al.* [19]. In particular, the results for M-M-$\mathcal{R}$ and M-M-Σ coincide as pointed out by Fofack *et al.*and, more generally, it is easy to see that G-M-$\mathcal{R}$ and G-M-Σ coincide in terms of the Laplace transforms of the stopped sum and the corresponding caching metrics (due to the TTLs' memorylessness property). An explicit distribution law for the inter-miss process in the case of the M-D-$\mathcal{R}$ model is open, albeit it can be obtained numerically.

We point out that these results hold for a single cache only. That means that the attractive addition property of the distribution law cannot be iterated (because the inter-miss process at a M-M-$\mathcal{R}$, M-M-Σ, and M-M-$\min(\Sigma,\mathcal{R})$ model is not Poisson).

[2]This result was previously obtained by Fofack *et al.* [19].

5. FEEDFORWARD CACHE NETWORKS

In this section we prove that MAPs are remarkably suitable to model inter-request processes in a feedforward cache network, to the point that the associated superposition and input-output operations are quite straightforward.

MAPs generalize Poisson processes by allowing the inter-arrival times to be dependent and also to belong to the broad class of phase-type (PH) distributions (to be defined later). MAPs have been motivated in particular by the need to mitigate the modelling restrictions imposed by the exponential distribution. From an analytical perspective, MAPs are quite attractive not only due to their versatility (MAPs are in fact dense in a large class of point processes, see Asmussen and Koole [5]), but also due to their tractability. Let us next give a common definition of MAPs.

DEFINITION 6 (MARKOV ARRIVAL PROCESS (MAP)).
A Markov arrival process is defined as a pair of matrices $(\mathbf{D}_0,\mathbf{D}_1)$ with equal dimensions, or as a joint Markov process $(J(t),N(t))$. The matrix $Q := \mathbf{D}_0+\mathbf{D}_1$ is the generator of a background Markov process $J(t)$. The matrix $\mathbf{D}_0$ is non-singular and a subintensity[3], and contains the rates of the so-called hidden transitions which govern the change of $J(t)$ only. In turn, the matrix $\mathbf{D}_1$ contains the (positive) rates of the so-called active transitions which govern the change of both $J(t)$ and a counting process $N(t)$, i.e., if $J(t^-)=i$ and a transition (i,j) from $\mathbf{D}_1$ occurs at time t, then $J(t)=j$ and $N(t)=N(t^-)+1$.

For the sake of familiarizing with MAPs, let us represent a two-state Markov Modulated Poisson Process (MMPP), described in terms of a background Markov process $J(t)$ with two states (see Figure 3); depending on the state, arrivals can occur (and contribute to a counting process $N(t)$) at rates λ_1 and λ_2.

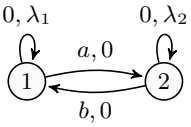

Figure 3: MAP representation of a MMPP; the transitions' components are hidden and active, respectively (e.g., in '$0,\lambda_1$', 0 is hidden and λ_1 is active)

The corresponding MAP is given by the hidden and active transition matrices

$$\mathbf{D}_0 = \begin{pmatrix} -a-\lambda_1 & a \\ b & -b-\lambda_2 \end{pmatrix},\ \mathbf{D}_1 = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix},$$

[3]A subintensity matrix S is similar to a stochastic matrix, except that rows sum to a non-positive value; formally, $S_{ii} < 0$, $S_{ij}\ge 0$ for $i\ne j$, and $\sum_{j=1}^m S_{ij}\le 0\ \forall i\in\{1,\dots,m\}$.

where $\mathbf{D}_0 + \mathbf{D}_1 = \begin{pmatrix} -a & a \\ b & -b \end{pmatrix}$ is the generator of $J(t)$.

In any state, $J(t)$ is driven by two competing transitions, i.e., a hidden one (either a or b) and an active one (either λ_1 or λ_2, respectively). One can note the two emerging features of the inter-arrival times of $N(t)$: they depend 1) on the state of the background Markov process $J(t)$, and 2) on the time to transiting $\mathbf{D}_0$ and to make a transition in $\mathbf{D}_1$.

5.1 MAPs for Two Simple Cache Networks

To introduce the main ideas of constructing MAPs for the input-output and superposition operations, we briefly present two simple examples of cache networks; the general results will be presented thereafter. For further more complex examples see Appendix B.

5.1.1 Input-Output

We first illustrate the input-output operation in a line-network scenario as in Section 4. Let the network from Figure 1 consist of two Σ caching nodes C_1 and C_2; requests arrive at C_1 as a Poisson process with rate λ, and the TTLs are $\sim exp(\mu)$ and $\sim exp(\nu)$. At node C_1, the arrivals can be represented as a Poisson process $N(t)$, which is itself an elementary single-state MAP M_1 defined in terms of

$$\mathbf{D}_0 = (-\lambda),\ \mathbf{D}_1 = (\lambda)\ ,$$

and a background Markov process with generator $\mathbf{Q} = (0)$. See Figure 1.(a) for its graphical representation.

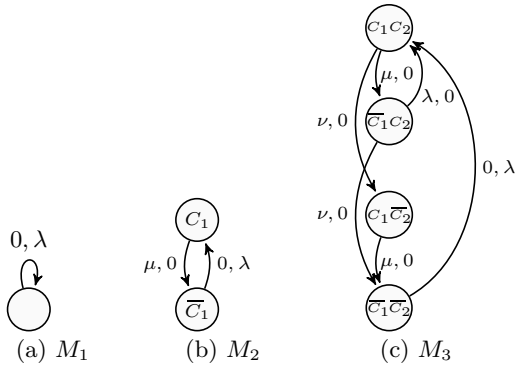

(a) M_1 (b) M_2 (c) M_3

Figure 4: M_1 corresponds to the arriving Poisson MAP at cache C_1 in Figure 1, M_2 to the output of C_1, and M_3 to the output of C_2

To construct the arrival MAP M_2 at C_2, capturing the inter-miss times at C_1, the basic idea is to duplicate the states of M_1 and suitably construct the hidden and active transitions. The new states (see Figure 1.(b)) are denoted by $\overline{C_1}$ (with the interpretation 'object is not in the cache') and C_1 (with the interpretation 'object is in the cache'). While in state $\overline{C_1}$, an arrival to C_1 triggers a miss—whence the active transition λ (i.e., the second component of '$0, \lambda$') to C_1. While in state C_1, the TTL may expire and hence the hidden transition μ to $\overline{C_1}$. It is important to remark that external requests while in C_1 result in hits and thus do not affect the MAP. Note also that the constructed MAP recovers that the inter-miss times (i.e., the time between two active transitions) are $hypo(\lambda, \mu)$ from Table 1. In matrix form, M_2 can also be represented as

$$\mathbf{D}_0' = \begin{pmatrix} -\lambda & 0 \\ \mu & -\mu \end{pmatrix} \text{ and } \mathbf{D}_1' = \begin{pmatrix} 0 & \lambda \\ 0 & 0 \end{pmatrix}\ .$$

Applying the same idea, we construct M_3 by duplicating the states of M_2. The four new states have the interpretations 'object is in none of the caches' (state $\overline{C_1}\overline{C_2}$), 'object is in only one cache' (states $C_1\overline{C_2}$ and $\overline{C_1}C_2$), and 'object is in both caches' (state C_1C_2). There are two important observations to make: one is that there is a single active transition (i.e., '$0, \lambda$') from $\overline{C_1}\overline{C_2}$ to C_1C_2. The other is that while in $\overline{C_1}C_2$, a request at C_1 does not result in an active transition because the object is already in C_2—and no miss at C_2 can occur—whence the *hidden* transition '$\lambda, 0$'. The remaining transitions are all hidden, capturing all possible TTLs' expirations depending on the states. In matrix form, M'' can also be represented as

$$\mathbf{D}_0'' = \begin{pmatrix} -\lambda & 0 & 0 & 0 \\ \mu & -\mu & 0 & 0 \\ \nu & 0 & -\lambda - \nu & \lambda \\ 0 & \nu & \mu & -\mu - \nu \end{pmatrix}\ ,$$

and $\mathbf{D}_1''$ contains only zeros except for λ on position $(1, 4)$.

5.1.2 Superposition

Consider now the tree topology from Figure 2. Applying the previous ideas, we can immediately construct the (independent) MAPs M_1 and M_2 corresponding to the inter-miss times at the caches C_1 and C_2, respectively (see Figures 5.(a)-(b)).

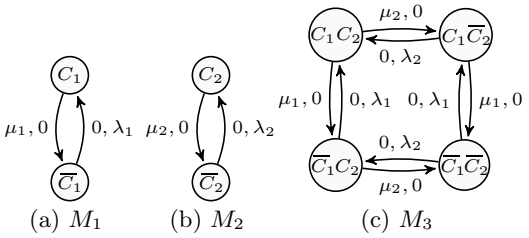

(a) M_1 (b) M_2 (c) M_3

Figure 5: The MAPs M_1, M_2, and M_3 corresponding to the inter-miss times at caches C_1 and C_2 from Figure 2, and their superposition

The construction of the superposition of M_1 and M_2, denoted by M_3, proceeds by forming the Cartesian product of the sets of states of M_1 and M_2; the resulting states have the same interpretation as in the previous subsection, e.g., $C_1\overline{C_2}$ stands for 'object in cache C_1 and not in cache C_2'. Moreover, the formation of the hidden and active transitions proceeds as before. For instance, while in state $C_1\overline{C_2}$ two transitions are possible: an active one (i.e., '$0, \lambda_2$') corresponding to an arrival at C_2, and a hidden one (i.e., '$\mu_1, 0$') corresponding to the TTL expiration.

Furthermore, one can construct the MAP corresponding to the inter-miss times at cache C_3 (in Figure 2) following the ideas so far. As the resulting number of states is eight (i.e., from doubling the states of M_3), we omit the graphical depictions. We can remark however that both the input-output and superposition operations result in an exponential increase of the number of MAP states; this fact will be elaborated more precisely later.

5.2 General Results

We now present the general results for constructing MAPs in feedforward networks of $\mathcal{R}$, Σ, and $\min(\Sigma, \mathcal{R})$ caches. As we previously mentioned, the MAP framework allows TTLs to belong to the broad class of PH distributions, which we define next.

DEFINITION 7 (PHASE-TYPE DISTRIBUTION). *Let* $\mathbf{S}$ *a* $m \times m$ *subintensity,* $\mathbf{S}_0 := -\mathbf{S}\mathbf{1}$*, and* π *a stochastic* m*-vector. Define a Markov process with generator*

$$\mathbf{P} := \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{S}_0 & \mathbf{S} \end{pmatrix} ,$$

which extends $\mathbf{S}$ *by an absorbing state* 0 *and exit transitions from every state in* $\mathbf{S}$ *to* 0*. A PH distribution (of order* m*), denoted as* $T = (\mathbf{S}, \pi)$*, is defined as the time until absorption in state* 0 *of the Markov process generated by* P*, and which starts in any of the states* $\{1, \ldots, m\}$ *according to* π*.*

We remark that we chose the less standard notation with the $\mathbf{0}$ vector on the first row instead of the last; this choice will permit expressing the input-output cache operation in a convenient manner.

Next we summarize the known result of MAPs' superposition and then present our main results on the input-output cache operation involving MAP requests and PH TTLs.

5.2.1 Superposition

First we briefly review the superposition of MAPs, for which we need to introduce the Kronecker sum $\oplus$ and product $\otimes$ operators for matrices.

If $\mathbf{A}$ and $\mathbf{B}$ are $m \times m$ and $n \times n$ matrices then

$$\begin{aligned} \mathbf{A} \otimes \mathbf{B} &:= \begin{pmatrix} a_{11}\mathbf{B} & \cdots & a_{1m}\mathbf{B} \\ \vdots & \ddots & \vdots \\ a_{n1}\mathbf{B} & \cdots & a_{mm}\mathbf{B} \end{pmatrix} \\ \mathbf{A} \oplus \mathbf{B} &:= \mathbf{A} \otimes \mathbf{I}_n + \mathbf{I}_m \otimes \mathbf{B} , \end{aligned}$$

where $\mathbf{A} = (a_{i,j})$ and $\mathbf{I}_k$ is the $k \times k$ identity matrix (note: the operator $\oplus$ is simplified for the case of square matrices).

THEOREM 2 (MAP SUPERPOSITION [29]). *If the MAPs* $M_1, \ldots, M_n$ *are represented in terms of the matrices* $(\mathbf{D}_0^1, \mathbf{D}_1^1)$, $\ldots, (\mathbf{D}_0^n, \mathbf{D}_1^n)$*, then their superposition* M *is also a MAP given by*

$$\begin{aligned} \mathbf{D}_0 &= \mathbf{D}_0^1 \oplus \cdots \oplus \mathbf{D}_0^n \\ \mathbf{D}_1 &= \mathbf{D}_1^1 \oplus \cdots \oplus \mathbf{D}_1^n . \end{aligned}$$

With abuse of notation we use the same operator $\oplus$ for the MAPs' superposition, i.e.,

$$M = M_1 \oplus \cdots \oplus M_n .$$

Consider for example $M_3 = M_1 \oplus M_2$ for the MAPs from Figure 5, and in particular the corresponding matrices of hidden transitions

$$D_0^1 = \begin{pmatrix} -\lambda_1 & 0 \\ \mu_1 & -\mu_1 \end{pmatrix}, \; D_0^2 = \begin{pmatrix} -\lambda_2 & 0 \\ \mu_2 & -\mu_2 \end{pmatrix} .$$

Then the Kronecker sum $D_0^1 \oplus D_0^2$ can be written as

$$\begin{pmatrix} -\lambda_1 & 0 & 0 & 0 \\ 0 & -\lambda_1 & 0 & 0 \\ \mu_1 & 0 & -\mu_1 & 0 \\ 0 & \mu_1 & 0 & -\mu_1 \end{pmatrix} + \begin{pmatrix} -\lambda_2 & 0 & 0 & 0 \\ \mu_2 & -\mu_2 & 0 & 0 \\ 0 & 0 & -\lambda_2 & 0 \\ 0 & 0 & \mu_2 & -\mu_2 \end{pmatrix} .$$

It is instructive to observe that the state-space of the Kronecker sum corresponds to the Cartesian product of the state spaces of M_1 and M_2 in lexicographical order, and which retains the Markovian properties of the (independent) superposed MAPs. Every state in M_1 corresponds to a block of 2 (i.e., the dimensionality of M_2) states in $M_1 \oplus M_2$; moreover, every state within such a block corresponds to a state in M_2. For the MAP M_3 from Figure 5.(c), the corresponding states are, in order, $\overline{C_1}\overline{C_2}$, $\overline{C_1}C_2$, $C_1\overline{C_2}$, and C_1C_2.

5.2.2 Input-Output: Σ, $\mathcal{R}$, and $\min(\Sigma, \mathcal{R})$ Caches

Let M be the MAP of cache requests and T be the TTL's PH distribution. We now prove that the input/miss process M', denoted formally using the notation

$$M' := M \oslash T ,$$

is also a MAP, for all Σ, $\mathcal{R}$, and $\min(\Sigma, \mathcal{R})$ caches. Note that, unlike in Section 4 where the $\mathcal{R}$ model was simpler than the Σ model, the opposite holds for MAPs for which reason we start with Σ.

THEOREM 3 (MAP-PH-Σ CACHE). *Consider a* Σ*-cache where requests arrive according to a MAP* $M = (\mathbf{D}_0, \mathbf{D}_1)$*. The TTLs are iid with a PH-distribution* T *and generator* $\mathbf{P}$*; also,* M *and* T *are independent. Then* $M' := M \oslash T$ *is a MAP with*

$$\begin{aligned} \mathbf{D}_0' &= (\mathbf{P} \oplus \mathbf{D}_0) + \begin{pmatrix} \mathbf{0} & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_1 & \ddots & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \ldots & \mathbf{0} & \mathbf{D}_1 \end{pmatrix} \\ \mathbf{D}_1' &= \begin{pmatrix} \mathbf{0} & \pi_1\mathbf{D}_1 & \pi_2\mathbf{D}_1 & \ldots & \pi_m\mathbf{D}_1 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \ldots & \mathbf{0} \\ \vdots & \vdots & \vdots & & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \ldots & \mathbf{0} \end{pmatrix} , \end{aligned}$$

where the $\mathbf{0}$ vectors have dimension $n \times n$. If M has n states and T has m transient and one absorbing states, then $\mathbf{D}_0'$ and $\mathbf{D}_1'$ are $n(m+1) \times n(m+1)$ matrices.

The proof is based on a constructive argument.

PROOF. First, it is easy to check that M' is a MAP according to Definition 6. The state space of M' is the Cartesian product of the state spaces of T and M (thus the term $\mathbf{P} \oplus \mathbf{D}_0$ in the expression of $\mathbf{D}_0'$). Note that, for technical reasons, the order of T and M in the Cartesian product is the opposite to the order in $M \oslash T$. The Cartesian product accounts for all the combinations of states from T and M. In particular, every state in T corresponds to a block of n states in M' (e.g., the first n rows and columns in $\mathbf{D}_0'$ and $\mathbf{D}_1'$), each corresponding to a state in M (recall the example after Theorem 2); moreover, block i corresponds to the states $(i-1)n + j \; \forall j = 1, \ldots, n$.

Next, to prove that M' models the miss process, we divide the $m+1$ blocks of M' into two groups: OUT and IN. The OUT group accounts for the situation when the object is 'out of the cache' and corresponds to the absorbing state of T, i.e., when the TTL is expired. While in any of the OUT states (corresponding to a position (i, j) in $\mathbf{D}_0'$ and $\mathbf{D}_1'$ with $1 \leq i \leq n$ and $1 \leq j \leq n(m+1)$), there are both hidden transitions (only due to the second Kronecker product in $\mathbf{P} \oplus \mathbf{D}_0$; the first product does not contribute because the current state of T is absorbing according to our representation of $\mathbf{P}$ from Definition 7) and active transitions (see the first block of rows in $\mathbf{D}_1'$). An active transition regenerates the phase of the TTL according to the stationary distribution π and consequently M' jumps to an IN block.

The IN group accounts for the situation when the object is 'in the cache', and each block within corresponds to one of the phases of T. While in any of the IN states (corresponding to a position (i,j) in $\mathbf{D}_0'$ and $\mathbf{D}_1'$ with $(n+1) \leq i \leq n(m+1)$ and $1 \leq j \leq n(m+1)$) there are only hidden transitions. Some are given by the entries of $\mathbf{P} \oplus \mathbf{D}_0$, and thus modelling the joint evolution of M and T. Importantly, we remark that since M' is within an IN group, the active transitions from $\mathbf{D}_1$ become passive; this is expressed in the second term of $\mathbf{D}_0'$. Moreover, the time between any two consecutive such transformed passive transitions corresponds to an element X_s from the definition of the stopping time of a Σ-cache (recall Eq. (3)). Finally, M' eventually jumps to the OUT block when an exit transition from $\mathbf{T}$ occurs.

Note that the proof implicitly uses the fact that the superposition of independent MAPs retains the underlying Markovian properties. □

Constructing the output for a $\mathcal{R}$ cache follows along the same lines except that the state of the TTL is reset with each arrival while the object is in the cache. This difference is modelled explicitly in the second term of D_0' in the following theorem.

THEOREM 4 (MAP-PH-$\mathcal{R}$ CACHE). *Under the same conditions as in Theorem 3, but for a $\mathcal{R}$ cache, $M' := M \oslash T$ is a MAP with*

$$\mathbf{D}_0' = (\mathbf{P} \oplus \mathbf{D}_0) + \begin{pmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \pi_1\mathbf{D}_1 & \pi_2\mathbf{D}_1 & \dots & \pi_m\mathbf{D}_1 \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \mathbf{0} & \pi_1\mathbf{D}_1 & \pi_2\mathbf{D}_1 & \dots & \pi_m\mathbf{D}_1 \end{pmatrix}$$

$$\mathbf{D}_1' = \begin{pmatrix} \mathbf{0} & \pi_1\mathbf{D}_1 & \pi_2\mathbf{D}_1 & \dots & \pi_m\mathbf{D}_1 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{pmatrix}.$$

where the $\mathbf{0}$ vectors have dimension $n \times n$. If M has n states and T has m transient and one absorbing states, then $\mathbf{D}_0'$ and $\mathbf{D}_1'$ are $n(m+1) \times n(m+1)$ matrices.

PROOF. The proof is identical to the previous one, except for accounting for the difference between Σ and $\mathcal{R}$ caches (see Eq. (3) vs. Eq. (2)). Concretely, while in the states of the IN group, an active transition becomes passive (as in the Σ case), but it also resets the phase of the TTL according to the probability vector π. □

Finally, the case of a $\min(\Sigma, \mathcal{R})$ cache exploits the known property that PH-distributions are closed under the minimum operator [10]):

LEMMA 3 (MINIMUM OF TWO PH-DISTRIBUTIONS). *Let $T_1 = (S_1, \pi_1)$ of order m, and $T_2 = (S_2, \pi_2)$ of order q be two PH distributions. Then $\min(T_1, T_2)$ is a PH distribution of order mq, and given by (S, π) where*

$$S = S_1 \oplus S_2 \text{ and } \pi = \pi_1 \otimes \pi_2 \ .$$

The construction of the $\min(\Sigma, \mathcal{R})$ output next follows by leveraging the property that the minimum of the two stopping times corresponding to Σ and $\mathcal{R}$, respectively, carries over to the TTLs' PH representations. In fact, given the minimum PH distribution from Lemma 3, the construction of the new matrix D_0' is comparable to the one from Theorem 4, and follows by repeating its construction for the second term of D_0'.

THEOREM 5 (MAP-PH-$\min(\Sigma, \mathcal{R})$ CACHE). *Consider a $\min(\Sigma, \mathcal{R})$ cache where requests arrive according to a MAP $M = (\mathbf{D}_0, \mathbf{D}_1)$. The two TTLs are iid with PH distributions T^Σ and $T^\mathcal{R}$ for Σ and $\mathcal{R}$, respectively. Arrivals and both TTLs are independent. If $\mathbf{P}$ denotes the generator of the PH distribution $\min(T^\Sigma, T^\mathcal{R})$ with corresponding initial vector $\pi = \pi^\Sigma \otimes \pi^\mathcal{R}$ (see Lemma 3, where π^Σ and $\pi^\mathcal{R}$ are the initial vectors of T^Σ and $T^\mathcal{R}$, respectively), then $M' := M \oslash \min(T^\Sigma, T^\mathcal{R})$ is a MAP with*

$$\mathbf{D}_0' = (\mathbf{P} \oplus \mathbf{D}_0) + \begin{pmatrix} \mathbf{0}_{n\times n} & \mathbf{0}_{n\times nq} & \mathbf{0}_{n\times nq} & \dots & \mathbf{0}_{n\times nq} \\ \mathbf{0}_{nq\times n} & \boldsymbol{\Omega} & \mathbf{0}_{nq\times nq} & \dots & \mathbf{0}_{nq\times nq} \\ & \mathbf{0}_{nq\times nq} & \boldsymbol{\Omega} & & \vdots \\ \vdots & \vdots & & \ddots & \mathbf{0}_{nq\times nq} \\ \mathbf{0}_{nq\times n} & \mathbf{0}_{nq\times nq} & & \dots & \boldsymbol{\Omega} \end{pmatrix}$$

$$\mathbf{D}_1' = \begin{pmatrix} \mathbf{0}_{n\times n} & \pi_1\mathbf{D}_1 & \pi_2\mathbf{D}_1 & \dots & \pi_{mq}\mathbf{D}_1 \\ \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \dots & \dots & \mathbf{0}_{n\times n} \\ \vdots & \vdots & & & \vdots \\ \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \dots & \dots & \mathbf{0}_{n\times n} \end{pmatrix},$$

where

$$\boldsymbol{\Omega} = \begin{pmatrix} \pi^\mathcal{R}{}_1\mathbf{D}_1 & \dots & \pi^\mathcal{R}{}_q\mathbf{D}_1 \\ \vdots & & \vdots \\ \pi^\mathcal{R}{}_1\mathbf{D}_1 & \dots & \pi^\mathcal{R}{}_q\mathbf{D}_1 \end{pmatrix}$$

has dimension $(nq \times nq)$, whereas the $\mathbf{0}_{a\times b}$ vectors have dimension $a \times b$. If M has n states, T^Σ has m transient states, and $T^\mathcal{R}$ has q transient states, then $\mathbf{D}_0'$ and $\mathbf{D}_1'$ are $n(m\,q+1) \times n(m\,q+1)$ matrices.

We make the important observation that because the order of T_1 and T_2 from Lemma 3 matters for the order of the states in the Markov chain of the corresponding minimum, the order of $\min(T^\Sigma, T^\mathcal{R})$ cannot be interchanged without changing the structure of D_0'.

PROOF. The eviction event for the minimum of the two stopping times as defined in Eq. (4) translates into either reaching the accepting state of T^Σ, or reaching the accepting state of $T^\mathcal{R}$ without an intermittent arrival. The minimum distribution of T^Σ and $T^\mathcal{R}$ captures this behavior up to the resetting of $T^\mathcal{R}$ upon arrivals. As mentioned in the proof of Theorem 4, an arrival resets $T^\mathcal{R}$ back to its initial state defined by its initial vector $\pi^\mathcal{R}$; however, the state of T^Σ is preserved. By Lemma 3 (and the underlying Kronecker sum), the states in $\min(\Sigma, \mathcal{R})$ are lexicographically ordered with the states of Σ followed by the states of $\mathcal{R}$. According to this order and because the reset behavior only changes the state of $T^\mathcal{R}$, an arrivals' effect remains local to each diagonal block Ω. Each Ω corresponds to the $\mathcal{R}$ reset matrix, as defined in the second term of D_0' in Theorem 4. □

In order to analyze feedforward cache networks we next complete the description of network cache operations by introducing the splitting concept of MAPs.

5.2.3 Probabilistic Splitting of Arrivals

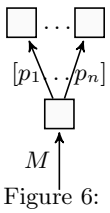

Figure 6: n-fold split of M

Apart from the superposition and the input-output operations, the third operation allows to split the MAP of an input (output) process, as shown in Figure 6. This works as an inverse operation of the superposition operator and allows to model the behavior of a cache feedforward network. A common splitting operation is when the process is split accordingly to some fixed probabilities. Such a construction allows to capture the behavior of an idealized load balancer.

LEMMA 4 (SPLITTING). *Assume a MAP $M = (D_0, D_1)$ is split into n sub processes according to a stochastic n-vector p. The resulting processes are characterized by the MAPs $M_i = (D_0^i, D_1^i)$, where*

$$D_0^i = D_0 + (1 - p_i)D_1 \text{ and } D_1^i = p_i D_1 \ ,$$

for $1 \leq i \leq n$.

PROOF. This is an extension of the known result that a single MAP is closed under thinning [33]. □

We point out that an input process represented as a MAP can be split in different ways of which many can be captured by a thinning operation. As further examples, the MAP arrival model is also closed under splitting requests according to their origin, or more generally, when the splitting decisions can be described by a Markov process. Besides accounting for splitting operations, our results can be further extended to account for various cache replication strategies as considered in Martina *et al.* [30].

Having introduced the main operations, we next state how to obtain the cache metrics from a MAP model of a cache.

5.3 Cache Metrics

For the MAP representation, the metrics can be derived in a uniform manner as the caching policies are encoded in the model.

LEMMA 5 (CACHE METRICS FOR A MAP CACHE). *For the $\mathcal{R}$, Σ, and $\min(\Sigma, \mathcal{R})$ caching policies, if the input process is an n-state MAP $M = (D_0, D_1)$, with steady-state probability vector p, and the output process is an n'-state MAP $M' = (D_0', D_1')$, with steady-state vector p', the miss probability and the cache occupancy are given by*

$$M = \frac{p' D_1' \mathbf{1}'}{p D_1 \mathbf{1}} \text{ and } H = 1 - M$$

$$\pi = \sum_{i=1}^{n'} p_i \, 1_{\{state \; i \; \in \; \mathsf{IN}\}}$$

where $\mathbf{1}$ and $\mathbf{1}'$ are all-ones vectors of dimensions $n \times 1$ and $n' \times 1$, respectively, and IN was defined in the proof of Theorem 3.

We note that the underlying stationary distribution of $M \oslash T$ exists when both M and T are ergodic.

5.4 Numerical Complexity

Although MAPs exactly characterize the miss-process of a cache network, the key drawback of the superposition and the input-output operations is that the state-spaces of the involved MAPs increase multiplicatively in the number of caches and the number of states of the TTLs' PH distribution, respectively (cf. Theorems 2-5).

We summarize the scaling behavior in the following lemma for the analysis of a binary tree.

LEMMA 6 (SCALING OF STATE SPACE). *Assume a complete binary tree of height h with 2^{h-1} arriving n-state MAPs. All nodes implement either $\mathcal{R}$ or Σ caches, with an m-state PH-distribution for the TTLs. Then, for a fixed object, the state space size for the exact analysis of the miss process scales as $n^{2^h} m^{2(2^h - 1)}$.*

The proof follows immediately by induction. We note that, in the case of $\min(\Sigma, \mathcal{R})$ caches, the state space complexity is even higher.

Although the computational complexity in the tree height appears to be prohibitively high, we point out that the matrices arising from the caching operations have a very regular sparse structure (as those involved in matrix geometric methods [32]). This sometimes allows to explicitly solve for the MAP's balance equations. Besides, matrix implementations using sparse representations enable the analysis of medium sized caching networks. In our own implementations, used as sanity checks for Theorems 3-5, we considered trees with height three (without relying on sparse-matrix representations) and trees up to height five (using an open-source sparse-matrix implementation), with multiple MAP arrivals and PH distributions.

6. CONCLUSION

In this paper we have provided the first exact analysis of TTL cache networks in great generality. We have developed two main methods covering three common TTL caching policies: $\mathcal{R}$, Σ, and $\min(\Sigma, \mathcal{R})$ which are employed in practical implementations. With the first method we have generalized an existing result available for lines of $\mathcal{R}$ caches with renewal requests, by additionally accounting for Σ and $\min(\Sigma, \mathcal{R})$. The key idea was to conveniently formalize the three TTL caching policies by a stopped-sum representation, whose transform could thereafter be computed using a change of measure technique. To address the lack of closure of renewals under superposition (and hence the inherent limitation to line networks), our second method proposed to use the versatile class of MAPs to model cache non-renewal requests. The key contribution was to show that MAPs are closed under the input-output operation of all three caching policies, whereby TTLs follow PH distributions. This property was instrumental for the exact analysis of feedforward networks. While the method addressing MAPs has a much broader applicability, it suffers however from an exponential increase in the space complexity. An immediate research direction concerns the development of fast numerical algorithms leveraging the sparse and regular structure of the underlying MAPs. In this way, one could gain analytical insight into the general cache placement problem in large heterogeneous networks, whereby different nodes can implement different policies.

7. REFERENCES

[1] Squid Web Cache FAQ. `http://wiki.squid-cache.org/SquidFaq/InnerWorkings`. accessed November 26, 2013.

[2] Amazon Web Service. *Amazon ElastiCache User Guide*, API version 2013-06-15 edition.

[3] S. Asmussen. Matrix-analytic models and their analysis. *Scandinavian Journal of Statistics*, 27(2):193–226, June 2000.

[4] S. Asmussen. *Applied probability and queues*, volume 2. Springer, 2003.

[5] S. Asmussen and G. Koole. Marked point processes as limits of Markovian arrival streams. *Journal of Applied Probability*, 30(2):365–372, June 1993.

[6] O. Bahat and A. M. Makowski. Measuring consistency in TTL-based caches. *Performance Evaluation*, 62(1):439–455, 2005.

[7] G. Bianchi, A. Detti, A. Caponi, and N. Blefari Melazzi. Check before storing: what is the performance price of content integrity verification in LRU caching? *ACM SIGCOMM Computer Communication Review*, 43(3):59–67, 2013.

[8] P. Billingsley. *Probability and Measure*. Wiley, 3 edition, 1995.

[9] S. Borst, V. Gupta, and A. Walid. Distributed caching algorithms for content distribution networks. In *Proceedings of IEEE INFOCOM*, pages 1–9, 2010.

[10] L. Breuer and D. Baum. *An introduction to queueing theory and matrix-analytic methods*. Springer, 2005.

[11] G. Casale, N. Mi, and E. Smirni. Bound analysis of closed queueing networks with workload burstiness. In *ACM SIGMETRICS*, pages 13–24, 2008.

[12] H. Che, Y. Tung, and Z. Wang. Hierarchical web caching systems: Modeling, design and experimental results. *IEEE Journal on Selected Areas in Communications*, 20(7):1305–1314, 2002.

[13] E. Cohen, E. Halperin, and H. Kaplan. Performance aspects of distributed caches using TTL-based consistency. In *Automata, Languages and Programming*, pages 744–756. Springer, 2001.

[14] E. Cohen and H. Kaplan. Aging through cascaded caches: Performance issues in the distribution of web content. In *Proceedings of ACM SIGCOMM*, pages 41–53, 2001.

[15] S. H. Cox, Y. Lin, and S. Wang. Multivariate exponential tilting and pricing implications for mortality securitization. *Journal of Risk and Insurance*, 73(4):719–736, 2006.

[16] A. Dan and D. Towsley. An approximate analysis of the LRU and FIFO buffer replacement schemes. In *Proceedings of ACM SIGMETRICS*, pages 143–152, 1990.

[17] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Springer, 2 edition, 1998.

[18] S. K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi, T. Koponen, B. Maggs, K. Ng, V. Sekar, and S. Shenker. Less pain, most of the gain: Incrementally deployable ICN. In *Proceedings of ACM SIGCOMM*, pages 147–158, 2013.

[19] N. C. Fofack, P. Nain, G. Neglia, and D. Towsley. Analysis of ttl-based cache networks. In *Proceedings of IEEE VALUETOOLS*, pages 1–10, 2012.

[20] C. Fricker, P. Robert, and J. Roberts. A versatile and accurate approximation for LRU cache performance. In *Proceedings of ITC*, pages 1–8, 2012.

[21] C. Fricker, P. Robert, J. Roberts, and N. Sbihi. Impact of traffic mix on caching performance in a content-centric network. In *IEEE NOMEN Workshop on Emerging Design Choices in Name-Oriented Networking*, pages 310–315, 2012.

[22] M. Gallo, B. Kauffmann, L. Muscariello, A. Simonian, and C. Tanguy. Performance evaluation of the random replacement policy for networks of caches. In *Proceedings of ACM SIGMETRICS/ PERFORMANCE*, pages 395–396, 2012.

[23] E. Gelenbe. A unified approach to the evaluation of a class of replacement algorithms. *IEEE Transactions on Computers*, 100(6):611–618, 1973.

[24] A. Gut. *Stopped Random Walks: Limit Theorems and Applications*. Springer, 2009.

[25] Y. T. Hou, J. Pan, B. Li, and S. S. Panwar. On expiration-based hierarchical caching systems. *IEEE Journal on Selected Areas in Communications*, 22(1):134–150, 2004.

[26] P. R. Jelenkovic. Asymptotic approximation of the move-to-front search cost distribution and least-recently used caching fault probabilities. *The Annals of Applied Probability*, 9(2):430–464, 1999.

[27] J. Jung, A. W. Berger, and H. Balakrishnan. Modeling TTL-based internet caches. In *Proceedings of IEEE INFOCOM*, pages 417–426, 2003.

[28] W. F. King III. Analysis of demand paging algorithms. In *IFIP Congress (1)*, pages 485–490, 1971.

[29] D. M. Lucantoni, K. S. Meier-Hellstern, and M. F. Neuts. A single-server queue with server vacations and a class of non-renewal arrival processes. *Advances in Applied Probability*, 22(3):676–705, Sept. 1990.

[30] V. Martina, M. Garetto, and E. Leonardi. A unified approach to the performance analysis of caching systems. In *Proceedings of IEEE INFOCOM (to appear)*, 2014. available: `http://arxiv.org/abs/1307.6702`.

[31] M. Musiela and M. Rutkowski. *Martingale methods in financial modelling*. Springer, 2005.

[32] M. F. Neuts. *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Dover, 1981.

[33] B. F. Nielsen. Note on the markovian arrival process. 1998.

[34] H. Pham. Some methods and applications of large deviations in finance and insurance. In *Paris-Princeton Lecture Notes in Mathematical Finance*. Springer, 2007.

[35] I. Psaras, R. G. Clegg, R. Landa, W. K. Chai, and G. Pavlou. Modelling and evaluation of CCN-caching trees. In *Proceedings of NETWORKING*, pages 78–91. Springer, 2011.

[36] L. Rizzo and L. Vicisano. Replacement policies for a proxy cache. *IEEE/ACM Transactions on Networking (ToN)*, 8(2):158–170, 2000.

[37] E. J. Rosensweig, J. F. Kurose, and D. F. Towsley. Approximate models for general cache networks. In *Proceedings of IEEE INFOCOM*, 2010.

APPENDIX

A. PROOFS OF SECTION 4

A.1 Proof of Proposition 1

Fix $t \geq 1$ and $x \geq 0$. The distribution $\tilde{F}(x)$ follows immediately from the integration rule from Eq. (10):

$$\tilde{\mathbb{P}}\left(X_t \leq x\right) = \int \frac{e^{-\omega X_t}}{\mathcal{L}_\omega(X)} 1_{\{X_t \leq x\}} d\mathbb{P}_{X_t} \ ,$$

where $\mathbb{P}_{X_t}$ is the projection of $\mathbb{P}$ on $\sigma(X_t)$. In turn, for $\tilde{T}(x)$, we have similarly

$$\begin{aligned}
\tilde{\mathbb{P}}\left(T_t \leq x\right) &= \int_{\Omega_{X_t} \times \{T_t \leq x\}} \frac{e^{-\omega X_t}}{\mathcal{L}_\omega(X)} d\mathbb{P}_{X_t} \times \mathbb{P}_{T_t} \\
&= \int_{T_t \leq x} \int_{\Omega_{X_t}} \frac{e^{-\omega X_t}}{\mathcal{L}_\omega(X)} d\mathbb{P}_{X_t} d\mathbb{P}_{T_t} \\
&= \int_{T_t \leq x} d\mathbb{P}_{T_t} = F_T(x) \ ,
\end{aligned}$$

where Ω_{X_t} denotes the (projected) sample space corresponding to X_t. In the first line we used the independence of X_t and T_t, i.e., the random vector (X_t, T_t) has the product measure $d\mathbb{P}_{X_t} \times \mathbb{P}_{T_t}$, where $\mathbb{P}_{X_t}$ and $\mathbb{P}_{T_t}$ are the projections of $\mathbb{P}$ on $\sigma(X_t)$ and $\sigma(T_t)$, respectively. In the second line we used Fubini's theorem.

Lastly, consider $B_1 \in \sigma(X_t)$ and $B_2 \in \sigma(T_t)$. Using again the independence of X_t and T_t (under $\mathbb{P}$) and Fubini's theorem we get

$$\begin{aligned}
\tilde{\mathbb{P}}\left(X_t \in B_1, T_t \in B_2\right) &= \int \frac{e^{-\omega X_t}}{\mathcal{L}_\omega(X)} 1_{\{X_t \in B_1, T_t \in B_2\}} d\mathbb{P} \\
&= \int_{\Omega_{X_t}} \frac{e^{-\omega X_t}}{\mathcal{L}_\omega(X)} 1_{\{X_t \in B_1\}} d\mathbb{P}_{X_t} \int_{T_t \in B_2} d\mathbb{P}_{T_t} \\
&= \tilde{\mathbb{P}}\left(X_t \in B_1\right) \tilde{\mathbb{P}}\left(T_t \in B_2\right) \ ,
\end{aligned}$$

which completes the proof. □

A.2 Proof of Corollary 1

Using the integration rule from Eq. (10) we first compute

$$\tilde{\mathbb{P}}(X \leq T) = \tilde{\mathbb{E}}\left[1_{\{X \leq T\}}\right] = \frac{\mathbb{E}\left[1_{\{X \leq T\}} e^{-\omega X}\right]}{\mathbb{E}\left[e^{-\omega X}\right]} = \frac{\psi(\omega)}{\mathcal{L}_\omega(X)} \ ,$$

such that the pmf of τ is

$$\begin{aligned}
\tilde{\mathbb{P}}(\tau = t) &= \left(\tilde{\mathbb{P}}(X \leq T)\right)^{t-1} \left(1 - \tilde{\mathbb{P}}(X \leq T)\right) \\
&= \left(\frac{\psi(\omega)}{\mathcal{L}_\omega(X)}\right)^{t-1} \left(1 - \frac{\psi(\omega)}{\mathcal{L}_\omega(X)}\right) .
\end{aligned}$$

Finally, applying Theorem 1 and manipulating progression series yields

$$\begin{aligned}
\mathbb{E}\left[e^{-\omega S_\tau}\right] &= \tilde{\mathbb{E}}\left[\mathcal{L}_\omega(X)^\tau\right] \\
&= \sum_{t=1}^{\infty} \left(\mathcal{L}_\omega(X)\right)^t \left(\frac{\psi(\omega)}{\mathcal{L}_\omega(X)}\right)^{t-1} \left(1 - \frac{\psi(\omega)}{\mathcal{L}_\omega(X)}\right) \\
&= \frac{\mathcal{L}_\omega(X) - \psi(\omega)}{1 - \psi(\omega)} \ ,
\end{aligned}$$

which completes the proof. □

A.3 Proof of Corollary 2

We first need to introduce the distribution convolution of S_t, for all $t \geq 1$, in the new space $(\Omega, \mathcal{F}, \tilde{\mathbb{P}})$. These are given for all $x \geq 0$ by by $\tilde{F}^1(x) := \tilde{F}(x)$ as in Proposition 1 and then recursively for $t > 1$ by the convolutions

$$\tilde{F}^t(x) = \int_0^x \tilde{F}^{t-1}(x-y) d\tilde{F}(y) \ ,$$

where $\tilde{F}^0(x) = 0$ for $x < 0$ and $\tilde{F}^0(x) = 1$ for $x \geq 0$. Assume also the existence of the corresponding densities $\tilde{f}^t$.

We next need the pmf of τ, which here proceeds by first conditioning on T, and then on $S_{t-1} = X_1 + \cdots + X_{t-1}$ and recalling from Proposition 1 that $\tilde{g}(x) = g(x)$.

$$\begin{aligned}
\tilde{\mathbb{P}}(\tau = t) &= \tilde{\mathbb{P}}(S_t > T,\, S_{t-1} \leq T) \\
&= \int_0^\infty \tilde{\mathbb{P}}(S_t > x,\, S_{t-1} \leq x)\, \tilde{g}(x)\, dx \\
&= \int_0^\infty \int_0^x (1 - \tilde{F}(x-y))\, \tilde{f}^{\,t-1}(y)\, dy\, g(x)\, dx \\
&= \int_0^\infty \left(\tilde{F}^{\,t-1}(x) \ - \ \tilde{F}^{\,t}(x)\right) g(x)\, dx \ .
\end{aligned}$$

The rest proceeds as in the proof of Corollary 1. □

A.4 Corollary for the G-G-min$(\mathcal{R}, \Sigma)$ Case

COROLLARY 3 (G-G-MIN$(\mathcal{R}, \Sigma)$).
Let τ be defined as in Eq. (4) and $G(\cdot)$, $g(\cdot)$ be the distribution and density of T^Σ, $H(\cdot)$ be the distribution of $T^\mathcal{R}$, and $F(\cdot)$ be the distribution of the arrivals. Then for some $\omega > 0$ the Laplace transform of the inter-miss time in the G-G-min$(\mathcal{R}, \Sigma)$ model is given by

$$\begin{aligned}
&\mathbb{E}\left[e^{-\omega\, S_\tau}\right] = \\
&\sum_{t \geq 1} \phi(\omega)^t \int_0^\infty \tilde{\mathbb{E}}\left[\prod_{i=1}^{t-1}(1 - H(X_i))\tilde{\mathbb{E}}\left[H(X_t)\right] | \sum_{i=1}^{t-1} X_i \leq k\right] \\
&- \tilde{\mathbb{E}}\left[\prod_{i=1}^{t-1}(1 - H(X_i))H(X_t) | \sum_{i=1}^{t} X_i \leq k\right] g(k)dk \\
&+ \int_0^\infty \tilde{\mathbb{E}}\left[\prod_{i=1}^{t-1}(1 - H(X_i))\, \tilde{\mathbb{E}}\left[H(X_t)\right] | \sum_{i=1}^{t-1} X_i \leq k\right] g(k)dk \\
&+ \int_0^\infty \tilde{\mathbb{E}}\left[\prod_{i=1}^{t-1}(1 - H(X_i))\tilde{\mathbb{E}}\left[1 - H(X_t)\right] | \sum_{i=1}^{t-1} X_i \leq k\right] \\
&- \tilde{\mathbb{E}}\left[\prod_{i=1}^{t}(1 - H(X_i)) | \sum_{i=1}^{t} X_i \leq k\right] g(k)dk
\end{aligned}$$

The proof is quite tedious and follows by conditioning; an explicit result for M-M-min$(\Sigma, \mathcal{R})$ is given in Table 1.

B. EXAMPLES FOR CACHES WITH PH TTLS

In this section we give examples for the application of Theorems 3-5. For all of the following three cache models (Σ, $\mathcal{R}$, min$(\Sigma, \mathcal{R})$), we assume MMPP arrivals, denoted by M, in the following reproduction of Figure 3:

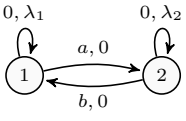

B.1 Σ Cache Model

We assume the following TTL T (in Markov chain representation and which starts in state 1 with probability one).

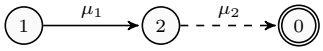

The output MAP is constructed by replicating the MMPP's states for each state of the TTL and adjusting for the inherent cache property, that no misses occur while the object is in the cache. This basic idea is reflected in taking the Cartesian product of T and M and subsequently making all of D_1's transitions passive (cf. definition of D_0' in Theorem 3).

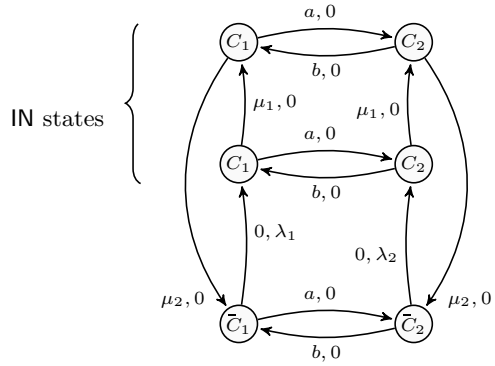

Note that we do not draw self-loops unless they are active transitions (i.e., the entries $D'_{0_{ii}}$ are not drawn, whereas an entry $D'_{1_{ii}}$ is drawn, as in the MMPP example). While the cache is in the IN state, further arrivals do not change the state of the cache. Thus, there are no transitions with λ_1 or λ_2 in the IN part of the resulting cache.

B.2 $\mathcal{R}$ Cache Model

Similarly to the Σ case we consider the MMPP arrival process M and the following TTL T (in Markov chain representation and which starts in state 1 with probability one):

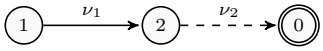

Constructing the output MAP for the $\mathcal{R}$ case bears a subtle difference from the Σ case. The basic idea is again to replicate the MMPP's states for each state of the TTL but then we have to accommodate for the $\mathcal{R}$ resetting behavior of this cache model: every arrival while the object is in the cache resets the TTL's state according to its initial vector.

Recalling the notations from Theorem 4, this idea is reflected by taking the Cartesian product of T and M and subsequently adjusting for the "resetting behavior", i.e., by making D_1's transitions passive and resetting T's state.

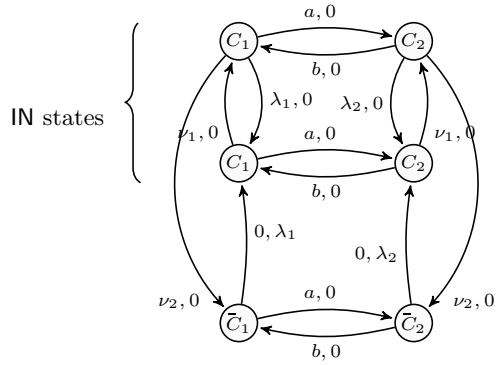

Note that if there was a greater number of TTL states and a non-trivial initial probability vector π for T, then the passive transitions '$\lambda_1, 0$' and '$\lambda_2, 0$' for each state of T would be directed to the initial states according to π and independently of T's current state. This is represented by the second term of D_0' in Theorem 4 by repeating the row with $\pi_i D_1$ for each state of the TTL.

Finally, we turn to the $\min(\Sigma, \mathcal{R})$ cache model which is more complicated due to the higher number of states involved.

B.3 $\min(\Sigma, \mathcal{R})$ Cache Model

Consider the same MMPP arrival process M with the following TTL representations. For the Σ part of the model, the TTL is called T^{Σ}:

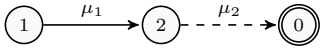

In turn, for the $\mathcal{R}$ part of the model, the TTL is called $T^{\mathcal{R}}$:

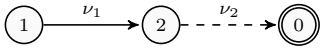

The corresponding output for a $\min(\Sigma, \mathcal{R})$ cache follows by constructing the PH minimum for $\min(T^{\Sigma}, T^{\mathcal{R}})$ which has four transient states and one absorbing state.

Then, we replicate the MMPP's states for each state of $\min(T^{\Sigma}, T^{\mathcal{R}})$ and link this Cartesian product construction with the reset behavior of the $\mathcal{R}$ model. As pointed out in the proof of Theorem 5, the resetting behavior of $\mathcal{R}$ has to preserve the state of T^{Σ}. This behavior is represented in the following output MAP by the two '$\lambda_1, 0$' and the two '$\lambda_2, 0$' transitions.

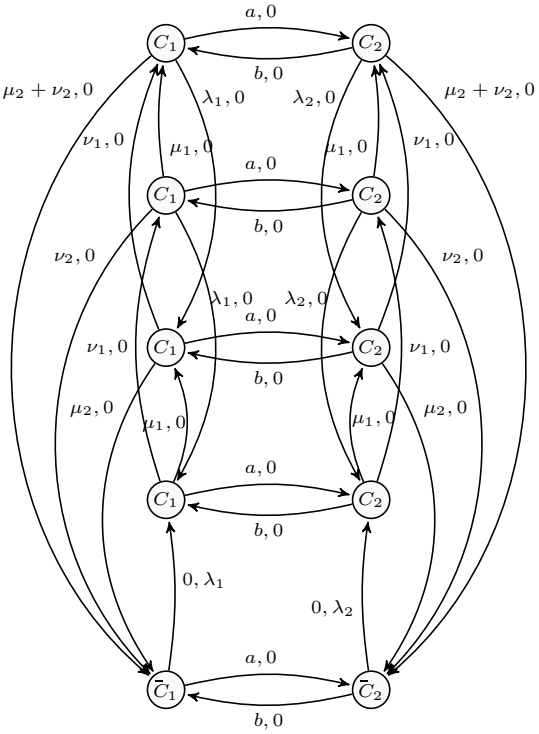

For a better visualization, the five layers of the above output MAP can be interpreted as 1) OUT (absorbing state), 2) IN $\Sigma-$State 1 $\mathcal{R}-$State 1, 3) IN $\Sigma-$State 2 $\mathcal{R}-$State 1, 4) IN $\Sigma-$State 1 $\mathcal{R}-$State 2, and 5) IN $\Sigma-$State 2 $\mathcal{R}-$State 2, from bottom up, where OUT and IN have the meanings from the proof of Theorem 3.